arXiv:1611.07065v1 [cs.NE] 21 Nov 2016

Recurrent Neural Networks With Limited Numerical Precision

Joachim Ott[*], **Zhouhan Lin**[‡], **Ying Zhang**[‡], **Shih-Chii Liu**[*], **Yoshua Bengio**[‡†]
[*]Institute of Neuroinformatics, University of Zurich and ETH Zurich
ottjoa+emdnn@gmail.com, shih@ini.ethz.ch
[‡]Département d'informatique et de recherche opérationnelle, Université de Montréal
[†]CIFAR Senior Fellow
{zhouhan.lin, ying.zhang}@umontreal.ca

Abstract

Recurrent Neural Networks (RNNs) produce state-of-art performance on many machine learning tasks but their demand on resources in terms of memory and computational power are often high. Therefore, there is a great interest in optimizing the computations performed with these models especially when considering development of specialized low-power hardware for deep networks. One way of reducing the computational needs is to limit the numerical precision of the network weights and biases, and this will be addressed for the case of RNNs. We present results from the use of different stochastic and deterministic reduced precision training methods applied to two major RNN types, which are then tested on three datasets. The results show that the stochastic and deterministic ternarization, pow2-ternarization, and exponential quantization methods gave rise to low-precision RNNs that produce similar and even higher accuracy on certain datasets, therefore providing a path towards training more efficient implementations of RNNs in specialized hardware.

1 Introduction

A Recurrent Neural Network (RNN) is a specific type of neural network which is able to process input and output sequences of variable length and therefore well suitable for sequence modeling. Various RNN architectures have been proposed in recent years based on different forms of non-linearity, such as the Gated Recurrent Unit (GRU) [3] and Long-Short Term Memory (LSTM) [9]. They have enabled new levels of performance in many tasks including speech recognition [1][2] and machine translation [7][4][22].

Compared to standard feed-forward networks, RNNs often take longer to train and are more demanding in memory and computational power. Thus, it is of great importance to accelerate computation and reduce training time of such networks.
Previous work showed the successful application of stochastic rounding strategies on feed forward networks, including binarization [6] and ternarization [13] of weights of vanilla Deep Neural Networks (DNNs) and Convolutional Neural Networks (CNNs) [18], and in [5] even the quantization of their activations during training and run-time. Quantization of RNN weights has so far only been used with pretrained models [19].

In this paper, a condensed and updated version of [16], we use different methods to reduce the numerical precision of weights in RNNs and test their performance on different benchmark datasets. We use two popular RNN models: vanilla RNNs, and GRUs. Section 2 covers the three ways of obtaining low-precision weights for the RNN models in this work, and Section 3 elaborates on the test results of the low-precision RNN models on different datasets. In [16] we additionally provide results and a possible explanation why binarization of RNNs does not work.

2 Rounding Network Weights

This work evaluates the use of 3 different rounding methods on the weights of two types of RNNs. These methods include the stochastic and deterministic ternarization method (TernaryConnect) [13], the pow2-ternarization method [21], and exponential quantization [16]. For all 3 methods, we keep a full-precision copy of the weights and biases during training to accumulate the small updates, while, during test time, we can either use the learned full-precision weights or use their deterministic low-precision version. As experimental results in Section 3 show, the network with learned full-precision weights usually yields better results than a baseline network with full precision during training, due to the extra regularization effect brought by the rounding method. The deterministic low-precision version could still yield comparable performance while drastically reducing computation and required memory storage at test time. We will briefly describe the former 3 low-precision methods, and elaborate on a additional method called Exponential Quantization we introduced in [16].

2.1 Ternarization and Pow2-Ternarization

TernaryConnect was first introduced in [13]. By limiting the weights to only 3 possible values, i.e., -1, 0 or 1 , this method does not require the use of multiplications. In the stochastic version of the method, the low-precision weights are obtained by stochastic sampling, while in the deterministic version, the weights are obtained by thresholding.

TernaryConnect allows weights to be set to zero. Formally, the stochastic form can be expressed as

$$W_{tern} = \text{sign}(W) \odot \text{binom}(p = \text{clip}(\text{abs}(2W), 0, 1)) \tag{1}$$

where $\odot$ is an element-wise multiplication, $\text{binom}(p)$ a function that draws samples from a binomial distribution, and $\text{clip}(x, a, b)$ clips values outside a given interval to the interval edges . In the deterministic form, the weights are quantized depending on 2 thresholds:

$$W_{tern} = \begin{cases} +1 & \text{if} \quad W > 0.5 \\ -1 & \text{if} \quad W \leq -0.5 \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

Pow2-ternarization is another fixed-point oriented rounding method introduced in [21]. The precision of fixed-point numbers is described by the Q*m.f* notation, where *m* denotes the number of integer bits including the sign bit, and *f* the number of fractional bits. For example, Q1.1 allows $(-0.5, 0, 0.5)$ as values. The rounding procedure works as follows: We first round the values to be in range allowed by the number of integer bits:

$$W_{\text{clip}} = \begin{cases} -2^m & \text{where} \quad W \leq -2^m \\ W & \text{where} \quad -2^m < W < 2^m \\ 2^m & \text{where} \quad W \geq 2^m \end{cases} \tag{3}$$

We subsequently round the fractional part of the values:

$$W_{p2t} = \text{round}(2^f \cdot W_{\text{clip}}) \cdot 2^{-f} \tag{4}$$

2.2 Exponential Quantization

Quantizing the weight values to an integer power of 2 is also a way of storing weights in low precision and eliminating multiplications. Since quantization does not require a hard clipping of weight values, it scales well with weight values.

Similar to the methods introduced above, we also have a deterministic and stochastic way of quantization. For the stochastic quantization, we sample the logarithm of weight value to be its nearest 2 integers, with the probability of getting one of them being proportional to the distance of the weight value from that integer. For weights with negative weight values we take the logarithm of its absolute vale, but add their sign back after quantization. i.e.:

$$log_2|W_b| = \begin{cases} \lceil log_2|W| \rceil & \text{with probability} \quad p = \frac{|W|}{2^{\lfloor log_2|W| \rfloor}} - 1 \\ \lfloor log_2|W| \rfloor & \text{with probability} \quad 1 - p \end{cases} \tag{5}$$

For the deterministic version, we set $log_2 W_b = \lceil log_2 W \rceil$ if the p in Eq. 5 is larger than 0.5.

Note that we just need to store the logarithm of quantized weight values. The actual instruction needed for multiplying a quantized number differs according to the numerical format. For fixed point representation, multiplying by a quantized value is equivalent to binary shifts, while for floating point representation it is equivalent to adding the quantized number's exponent to the exponent of the floating point number. In either case, no complex operation like multiplication is needed.

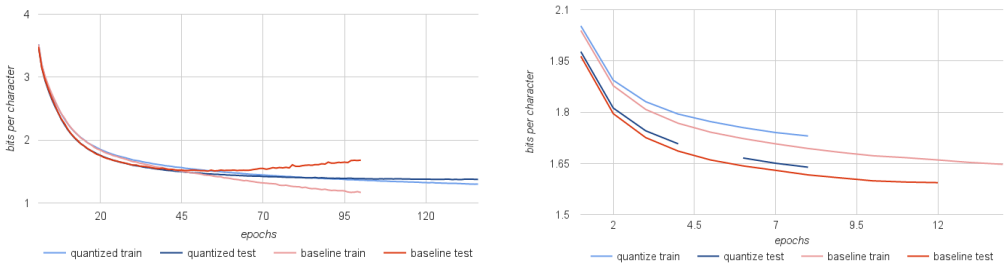

(a) (b)

Figure 1: Training and test set performance of vanilla RNNs learning on PTB (a) and text8 (b) datasets.

3 Experimental Results and Discussion

In the following experiments, we test the effectiveness of the different rounding methods on two different types of applications: character-level language modeling and speech recognition.

3.1 Vanilla RNN

We validate the low-precision vanilla RNN on 2 datasets: text8 and Penn Treebank Corpus (PTB).

The **text8** dataset contains the first 100M characters from Wikipedia, excluding all punctuations. It does not discriminate between cases, so its alphabet has only 27 different characters: the 26 English characters and space. We take the first 90M characters as training set and split them equally into sequences with 50 character length each. The last 10M characters are split equally to form validation and test sets.

The **Penn Treebank Corpus** [23] (PTC) contains 50 different characters, including English characters, numbers, and punctuations. We follow the settings in [14] to split our dataset, i.e., 5017k characters for training set, as well as 393k and 442k characters for validation and test set, respectively.

Model and Training The models are built to predict the next character given the previous ones, and performances are evaluated with the bits-per-character (BPC) metric, which is $\log_2$ of the perplexity, or the per-character log-likelihood (base 2). We use a RNN with ReLU activation and 2048 hidden units. We initialize hidden-to-hidden weights as identity matrices, while input-to-hidden and hidden-to-output matrices are initialized with uniform noise.

We can see the regularization effect of stochastic quantization from the results of the two datasets. In the PTB dataset, where the model size slightly overfits the dataset, the low-precision model trained with stochastic quantization yields a test set performance of 1.372 BPC that surpasses its full precision baseline (1.505 BPC) by around 0.133 BPC (Fig. 1, left). From the Figure 1 we can see that stochastic quantization does not significantly hurt training speed and manages to get better generalization when the baseline model begins to overfit. On the other hand, we can also see, from the results on the text8 dataset where the same sized model now underfits that the low-precision model performs worse (1.639 BPC) than its baseline (1.588 BPC). (Fig. 1b and Table 1).

Table 1: Results from Vanilla RNNs (VRNNs) with exponential quantization.

Dataset	RNN Type	Baseline	Exponential Quantization
text8	VRNN	1.588 BPC	1.639 BPC
PTC	VRNN	1.505 BPC	1.372 BPC

3.2 GRU RNNs

This section presents results from the various methods to limit numerical precision in the weights and biases of GRU RNNs which are then tested on the TIDIGITS dataset.

Dataset TIDIGITS[12] is a speech dataset consisting of clean speech of spoken numbers from 326 speakers. We only use single digit samples (zero to nine) in our experiments, giving us 2464 training

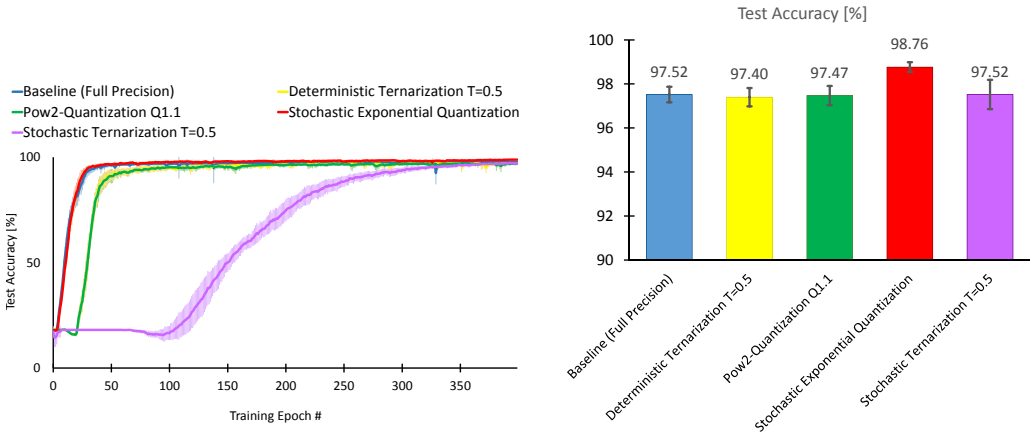

(a) (b)

Figure 2: GRU models trained on the TIDIGITS speech recognition task. (a) Effect of rounding methods applied on GRU weights and biases during training compared against baseline (blue). Thick curves show the mean of 10 runs, half-transparent areas around a curve show the variance. (b) Mean final test accuracy of 10 runs per method. Exponential quantization applied on GRU weights and biases leads to higher than baseline final accuracy.

samples and 2486 validation samples. The labels for the spoken 'zero' and 'O' are combined into one label, hence we have 10 possible labels. We create MFCCs from the raw waveform and do leading zero padding to get samples of matrix size 39x200. The MFCC data is further whitened before use.

Model and Training The model has a 200 unit GRU layer followed by a 200 unit fully-connected ReLU layer. The output is a 10 unit softmax layer. Weights are initialized using the Glorot & Bengio method [8]. The network is trained using Adam [11], and BatchNorm [10] is not used. We train our model for up to 400 epochs with a patience setting of 100 (no early stopping before epoch 100). GRU results in Figure 2 are from 10 experiments, with each experiment starting with a different random seed. This is done because previous experiments have shown that different random seeds can lead up to a few percent difference in final accuracy. We show average and maximum achieved performance on the validation set.

Effect of Rounding Methods To assess the impact of weight quantization, we trained our model and applied the rounding methods to all GRU weights and biases during training. Figure 2 (a) shows how rounding applied to GRU weights and biases has an effect on convergence compared to the full-precision baseline. Figure 2 (b) shows the mean final accuracy compared to baseline. These results show that limiting the numerical precision of GRU weights and biases using exponential quantization is beneficial for learning in this setup: the mean final accuracy is higher than baseline, while the convergence is as fast as baseline.

4 Conclusion and Outlook

This paper shows how low-precision quantization of weights can be performed effectively for RNNs. We presented 3 existing methods of limiting the numerical precision, and used them on two major RNN types and determined how the limited numerical precision affects network performance across 3 datasets. In the language modeling task, the low-precision model surpasses its full-precision baseline by a large gap (0.133 BPC) on the PTB dataset. We also show that the model will work better if put in a slightly overfitting setting, so that the regularization effect of stochastic quantization will begin to function. In the speech recognition task, we show that with stochastic ternarization we can achieve baseline results, and with exponential quantization even surpass the full-precision baseline. The successful outcome of these experiments means that lower resource requirements are needed for custom implementations of RNN models.

5 Acknowledgments

We are grateful to INI members Danny Neil, Stefan Braun, and Enea Ceolini, and MILA members Philemon Brakel, Mohammad Pezeshki, and Matthieu Courbariaux, for useful discussions and help with data preparation. We thank the developers of Theano[24], Lasagne, Keras, Blocks[25], and Kaldi[17].
The authors acknowledge partial funding from the Samsung Advanced Institute of Technology, University of Zurich, NSERC, CIFAR and Canada Research Chairs.

References

[1] Dario Amodei, Rishita Anubhai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Jingdong Chen, Mike Chrzanowski, Adam Coates, Greg Diamos, et al. Deep-Speech 2: End-to-end speech recognition in English and Mandarin. *arXiv*, page 28, 2015. URL `http://arxiv.org/abs/1512.02595`.

[2] William Chan, Navdeep Jaitly, Quoc V. Le, and Oriol Vinyals. Listen, attend and spell. *arXiv preprint*, pages 1–16, 2015. URL `http://arxiv.org/abs/1508.01211`.

[3] Kyunghyun Cho, Bart van Merrienboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv*, 2014. URL `http://arxiv.org/abs/1409.1259`.

[4] Junyoung Chung, Kyunghyun Cho, and Yoshua Bengio. A character-level decoder without explicit segmentation for neural machine translation. *arXiv*, 2016. URL `http://arxiv.org/abs/1603.06147`.

[5] Matthieu Courbariaux and Yoshua Bengio. BinaryNet: Training deep neural networks with weights and activations constrained to +1 or -1. *arXiv*, 2016. URL `http://arxiv.org/abs/1602.02830`.

[6] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. BinaryConnect: Training Deep Neural Networks with binary weights during propagations. *arXiv*, pages 1–9, 2015. URL `http://arxiv.org/abs/1511.00363`.

[7] Jacob Devlin, Rabih Zbib, Zhongqiang Huang, Thomas Lamar, Richard Schwartz, and John Makhoul. Fast and robust neural network joint models for statistical machine translation. *ACL*, 17:1370–1380, 2014. URL `http://acl2014.org/acl2014/P14-1/pdf/P14-1129.pdf`.

[8] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. *Proceedings of the 13th AISTATS*, 9:249–256, 2010. ISSN 15324435. doi: 10.1.1.207.2059. URL `http://machinelearning.wustl.edu/mlpapers/paper{_}files/AISTATS2010{_}GlorotB10.pdf`.

[9] Sepp Hochreiter, S Hochreiter, Jürgen Schmidhuber, and J Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–80, 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL `http://www.ncbi.nlm.nih.gov/pubmed/9377276`.

[10] Sergey Ioffe and Christian Szegedy. Batch Normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv*, 2015. ISSN 0717-6163. doi: 10.1007/s13398-014-0173-7.2. URL `http://arxiv.org/abs/1502.03167`.

[11] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv:1412.6980 [cs]*, pages 1–15, 2014. URL `http://arxiv.org/abs/1412.6980$\delimiter"026E30F$nhttp://www.arxiv.org/pdf/1412.6980.pdf`.

[12] R G Leonard and G Doddington. Tidigits Ldc93S10, 1993.

[13] Zhouhan Lin, Matthieu Courbariaux, Roland Memisevic, and Yoshua Bengio. Neural Networks with Few Multiplications. pages 1–8, 2015. URL `http://arxiv.org/abs/1510.03009`.

[14] Tomáš Mikolov, Ilya Sutskever, Anoop Deoras, Hai-Son Le, Stefan Kombrink, and J Cernocky. Subword language modeling with neural networks. *preprint*, 2012. URL `http://www.fit.vutbr.cz/imikolov/rnnlm/char.pdf`.

[15] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015. ISSN 0028-0836. doi: 10.1038/nature14236. URL `http://dx.doi.org/10.1038/nature14236`.

[16] Joachim Ott, Zhouhan Lin, Ying Zhang, Shih-Chii Liu, and Yoshua Bengio. Recurrent neural networks with limited numerical precision. *arXiv preprint arXiv:1608.06902*, 2016.

[17] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, Jan Silovsky, Georg Stemmer, and Karel Vesely. The Kaldi Speech Recognition Toolkit, 2011. URL `http://kaldi-asr.org`.

[18] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. XNOR-Net: ImageNet classification using binary convolutional neural networks. *arXiv preprint*, pages 1–17, 2016. URL `http://arxiv.org/abs/1603.05279`.

[19] Sungho Shin, Kyuyeon Hwang, and Wonyong Sung. Fixed-point performance analysis of recurrent neural networks. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 976–980. IEEE, 2016.

[20] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016. ISSN 0028-0836. doi: 10.1038/nature16961. URL `http://www.nature.com/doifinder/10.1038/nature16961`.

[21] Evangelos Stromatias, Daniel Neil, Michael Pfeiffer, Francesco Galluppi, Steve B Furber, and Shih-Chii Liu. Robustness of spiking Deep Belief Networks to noise and reduced bit precision of neuro-inspired hardware platforms . *Frontiers in Neuroscience*, 9(July):1–14, 2015. ISSN 1662-453X. doi: 10.3389/fnins.2015.00222. URL `http://www.frontiersin.org/Journal/Abstract.aspx?s=755{&}name=neuromorphic{_}engineering{&}ART{_}DOI=10.3389/fnins.2015.00222`.

[22] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. *Advances in Neural Information Processing Systems (NIPS)*, pages 3104–3112, 2014. URL `http://papers.nips.cc/paper/5346-sequence-to-sequence-learning-with-neural`.

[23] Ann Taylor, Mitchell Marcus, and Beatrice Santorini. the Penn Treebank: an Overview. *Treebanks*, pages 5–22, 2003. doi: 10.1007/978-94-010-0201-1_1. URL `http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.9.8216{&}rep=rep1{&}type=pdf`.

[24] Theano Development Team. Theano: A {Python} framework for fast computation of mathematical expressions. *arXiv e-prints*, abs/1605.0, 2016. URL `http://arxiv.org/abs/1605.02688`.

[25] Bart van Merrienboer, Dzmitry Bahdanau, Vincent Dumoulin, Dmitriy Serdyuk, David Warde-farley, Jan Chorowski, and Yoshua Bengio. Blocks and Fuel : Frameworks for deep learning. pages 1–5, 2015. URL `http://arxiv.org/abs/1506.00619`.